

## **A Design Strategy for Autonomous Systems**

**Pete Forster**

**Department of Artificial Intelligence  
Edinburgh University  
5 Forrest Hill  
Edinburgh EH1 2QL  
SCOTLAND**

### **Abstract**

This paper identifies and puts forward some solutions to crucial issues regarding the competent performance of an autonomously operating robot, namely that of handling multiple and variable data sources containing overlapping information and maintaining coherent operation whilst responding adequately to changes in the environment. Support for the ideas developed for the construction of such behaviour are extracted from speculations in the study of cognitive psychology, an understanding of the behaviour of controlled mechanisms and the development of 'behaviour-based' robots in a few robot research laboratories. The validity of the ideas expressed is supported by some simple simulation experiments in the field of mobile robot navigation and guidance.

# 1 Introduction

In many robot application scenarios, a human is in close communication with the robot. The human may be closely involved in the execution of tasks or merely have a supervisory role. In either case, malfunction or erroneous operation by the robot can be rapidly observed and hopefully prevented or corrected by the operator. In more remote environments (eg. space and subsea), it may not be possible to utilize human involvement to such an extent due to the limitations of communication. Thus robots operating in such environment must be able to operate autonomously.

What properties does such a robot need to possess and how can they be achieved? The answer to both of these questions is not obvious. The issue of autonomy has already been addressed, but it is not sufficient merely to state that autonomy is the basic requirement for independently operating robots. Future generations of robots will continue to interact with humans, if only to be switched on and off. Robots are effectively the slaves of humans and must be able to receive and react to instructions from them. From this point, the term "autonomous robot" should be taken to mean a robot that is not only able to operate in some autonomous manner, but is also capable of benefiting from information or directives contributed by humans or other sources. Communication channels must not only exist between human and robot, but also between robots and between a robot and any other information sources. Thus an autonomous robot might be summarized as shown in Figure 1.




Figure 1: Input-Output schematic for an autonomous robot

Not all inputs to the robot are controllable. The robot may be continually perturbed by the environment; for example, an underwater robot may be subjected to drift by currents and a land-based rover is subject to being rocked or skidding. These inputs are categorized as 'disturbances'. Physical response includes both the response of the robot and its effectors.

## 2 Some Specifications

Specifying the necessary behaviour for an autonomous robot to operate competently is not only highly domain dependent, but can result in very lengthy top-down breakdowns. This in itself may influence the way a robot's systems are architected. This paper takes a different angle by considering in more general terms the necessary characteristics of an intelligently acting robot. Intelligence in this context means 'competence in the robot's operating domain'. The criteria under consideration are:

- the robot must be able to operate autonomously;
- the robot should be capable of flexibly utilizing any obtained knowledge which could benefit its performance;
- the robot must be able to respond adequately fast to any scenario that it may face whilst maintaining coherent task strategies;
- the robot should be capable of communicating with and accomodating communications from external sources or agents;

The topic of autonomous operation has already been raised which prompted the issue of how to orchestrate and process inputs from the robot's sensors and other sources with any communication from human sources. Thus the second criteria is specified so that the robot is able to respond as effectively as possible given the available data. This concept allows for degradation and failing of inputs to the robot. It is envisaged that considerable redundancy or overlapping of sensor data will be available to the robot.

A major problem incurred when using conventional sequential computing hardware is that only one process can be executing at a given instant. In real-time control and simulation systems, a scheduler might be used to run different processes on different processing cycles so that a given process is repeated regularly enough to satisfy the update rate requirement of the associated control or simulation process. There is nevertheless a computational limitation to this approach. What is required is specific hardware to run many processes sufficiently fast to meet the needs of the whole system. For example, if a subsea robot had reached a target site and was engaged in task planning activity, it must still maintain vigilance within the environment to avoid falling debris, say, or merely hold its position against variable currents. There should be no reason for its reaction time to deteriorate because it is engaged in simultaneous activity. However, intelligent behaviour is not the result of a package of independently running processes. These processes must not only be configured in such a way that necessary response times can be achieved, but also so that a coherent behaviour emerges for the whole system.

### 3 System Design

There are two areas of study which contribute directly to the realisation of systems which meet the design criteria specified: the study of animate behaviour and the study of mechanism control.

#### 3.1 Animate behaviour

There is no shortage of autonomously behaving creatures to study in order to gain an understanding of how a machine might be constructed so as to exhibit intelligent behaviour. However, the topic is still surrounded by controversy and leading researchers can still only speculate in the main on how animate behaviour arises. Nevertheless, schemas developed in order to assist us to comprehend the function of biological processing mechanisms can provide useful principles for robot design techniques. Stillings [Stillings et al, 1988] describes human skill behaviour as arising from a combination of *controlled* and *automatic* processes. During the skill learning phase, he interprets subsequent improvement as being a transition of process bias from the slow and methodical controlled process domain to the rapid automatic process domain. He further identifies our limitations in only being able to cope with one 'high-level' controlled process at a time, and the limited control we have over our memory. Computational machines can be designed to overcome these limitations.

Robot processing is conducted on non-biological hardware which suggests that intelligent behaviour should not necessarily be constructed by replicating biological processing, but by determining methods that exploit the characteristics of this more understandable device. There are clearly areas in which processing strategies developed from neurological studies might be appropriate (eg. vision), but believing that all processing could be conducted using similar techniques is optimistic, if not naive.

#### 3.2 Mechanism behaviour

An alternative domain in which to extract concepts for the design of intelligent behaviour is that of mechanisms, particularly those using feedback control processes. Figure 2 outlines phases involved in the response of a system to inputs. The physical system consists of the basic structure of the machine together with the necessary effectors (output mechanisms) with which the machine response may be controlled (eg. steering gear). The physical system is subject to being disturbed by external physical influences such as gusts, rough terrain etc.

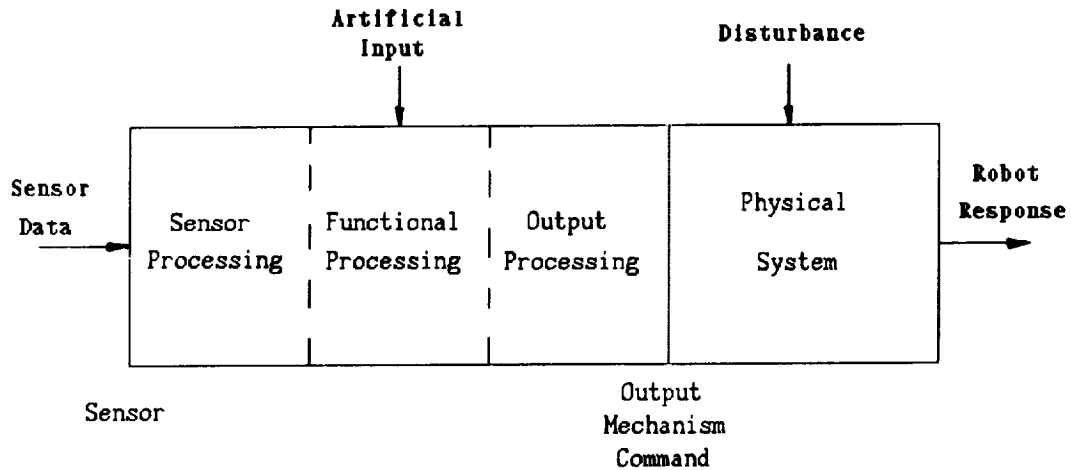


Figure 2: Breakdown of a controlled system

Processing to control the machine is usually conducted on processors, but can be mechanical (eg. directional control of a wind-mill). The processing element issues commands to the effector control devices which modify the response of the whole machine. This processing can be split into three components for most, if not all, controlled mechanism designs. Initially, raw sensor data is processed into a more usable form. This sensor system may govern the way the sensed data is used (as in the *collision avoidance* process described later). This refined data is then used by the functionally oriented processing to produce response demands which can be translated by further output processing into effector demands.

Mechanism processing does not only have to control effectors that directly generate machine response, but they may also have to control the hardware containing the sensors (eg. control of a camera position, orientation, focus and zoom). By controlling the data supplied to the sensor system in this way, this output can be used to indirectly drive the machine response as shown in Figure 3.

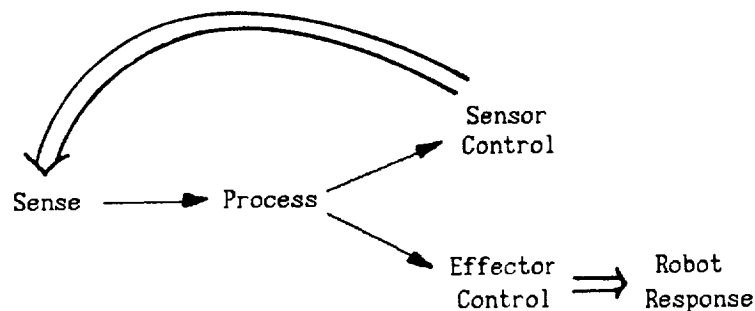


Figure 3: System control paths

### 3.3 Behaviour-based systems

Several prominent AI researchers have recognized the need to design autonomous systems using *behaviours* as a building block. Brooks [Brooks, 1987] at M.I.T. has designed autonomous mobile robots in this way using a *subsumption* architecture and robust *behaviours* in a distributed processing implementation. Raibert [Raibert, 1986] has successfully built multi-legged robots by commencing his design with a single leg implementation before progressing to the multi-legged case. These approaches attempt to obtain fast and competent response to the environment before incorporating any explicit environment interpretation and reasoning. By orchestrating these reactive behaviours in a coherent manner, more sophisticated competent behaviour may emerge. The problems incurred in the design process are that of specifying what behaviours should be explicitly designed and how they might be orchestrated to construct an intelligent robot or system.

### 3.4 A processing configuration

Figure 4 describes a processing arrangement to implement an intelligent processing system based on the concepts expressed in the previous three sections. The architecture attempts to ensure that the system responds adequately fast to the environment by executing multiple processes in a distributed fashion. Automatic or subcognitive behaviour lies at the lower end of the diagram where the computation time is short for simple sensors, equivalent to proprioception in animate behaviour, to initiate a response. Types of behaviour found in this category would be collision or hazard avoidance and general reflexive actions. Feeding into these low-level processes are the higher level controlled or cognitive behaviour processes.

Several methods of feeding in to the low-level processes without significantly increasing response times have been investigated. The first method consists of directly combining output response demands (by simple summation or by taking minimums or maximums) after the functional processing stage has been completed. Other methods consist of interfering with the sensor system processing to partially control the behaviour of the low-level processes. This can be performed either by suppressing/enhancing the sensitivity of components of the sensor system, or causing 'hallucinations' in the sensor system to evoke particular response. Both of these methods are examples of how a higher-level can exploit the characteristics of lower level processes. Particular examples are described in the example in the following section.

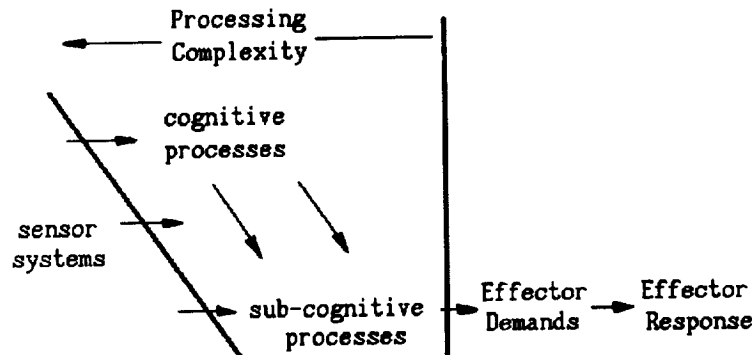


Figure 4: Processing configuration

### 3.5 A flexible architecture

An autonomous robot is likely to possess a multitude of sensing devices. Each device will provide a unique set of data but the contents of which are likely to be duplicated or part-duplicated by other sensor devices. The quality of the data returned by sensors is likely to vary according to the environment and those sensors may degrade or fail during the operation of the robot. Thus robots should be designed to flexibly utilize any information sources available to them to be able to perform as well as possible within the constraints imposed on them. Methods of amalgamating such sensor derived data have been established [Durrant-Whyte, 1987] [Forster et al, 1988] by manipulating multi-sensor state and estimated variance data. Hallam [Hallam, 1985] has developed self-navigation algorithms which are able to flexibly accommodate various types of data in deriving navigational information. This is described in more detail at the end of the following section.

## 4 Motion Guidance Example

The field of motion guidance is a prime application area both due to the multiple considerations present and the considerable amount of work already performed in this domain. A *collision avoidance* algorithm has been designed and implemented in a computer simulation for an autonomous vehicle. A simple sensor system consisting of a range finding device of limited azimuth resolution is used to achieve the avoidance behaviour. The algorithm processes data received from this sensor system and produces speed and turn-rate demands. In a wheeled vehicle, this would require the drive power/gearing and steering to be controlled. These are the 'effectors' that will generate vehicle response. The algorithm is notable in that it does not use any memory storage between processing cycles and does not involve any explicit interpretation of the sensor data to produce a map of the sensed environment. Such activity complicates the processing task which leads to an increasing response time. Experiments have shown that the system is capable of exhibiting quite complex motion behaviour. Figure 5 illustrates typical behaviour.




Figure 5: Mobile robot collision avoidance behaviour

The vehicle trail is shown as a dotted line, increasing density indicating a reduction in speed. The sensor system is represented by a number of whiskers protruding from the vehicle. If these whiskers intersect a wall, then the proportion of whisker length obstructed is processed by the algorithm to generate avoidance motion. The number of whiskers and their azimuth spread increases as the sensed environment becomes increasingly cluttered. This will also initiate a reduction in vehicle speed which enables it to perform smaller radius turns.



A simple direction seeking algorithm, based on proportional navigation techniques, has been designed to demonstrate *location seeking* behaviour. This algorithm processes goal direction data and, as for the *collision avoidance* behaviour, produces speed and turn-rate demands. The behaviour consists of the vehicle aiming towards the goal direction. Figure 6 illustrates an example of this behaviour in free space with the vehicle starting from rest, facing in a direction almost 180 degrees away from the goal direction. No range information is provided, so this algorithm results in the vehicle overshooting the goal.




Figure 6: Mobile robot location seeking behaviour

The *location seeking* process is not independent of other processes since it requires direction data. This could be supplied directly from a sensor system (eg. visual) if the location could be sensed. However, a 'higher level' process could formulate a plan with suitable way-points and feed these to the location seeker; this is an example of how a high level process can exploit the capability of a low level process.

The collision avoidance and location seeking processes have potentially conflicting effector output demands. If a resultant competent behaviour is to be achieved, the two processes must be combined so that a single pair of effector output demands are derived. Without further processing of sensed data, the collision avoidance behaviour should maintain dominance over the location seeking behaviour but the resulting behaviour should be seen to take the demands of the location seeker into account whilst the collision avoidance behaviour is stimulated. Changing between processes does not produce a competent behaviour. A successful method of amalgamation has been achieved by modifying the *collision avoidance* algorithm so that its outputs were unbounded and by creating a merging algorithm to combine the two pairs of outputs and bound the result. The merger algorithm simply takes the minimum of the speed demands and the sum of the turn rate demands. This combined structure generates *trajectory guidance* behaviour. An example validation of this behaviour is shown in Figure 7.




Figure 7: Mobile robot combined motion guidance behaviour

The above two examples of behaviour-based processes are just potential components of a complete autonomous vehicle. More competent behaviour could include the ability to perform navigation and guidance to achieve higher level objectives or tasks. Meeting this particular strategy, a self-navigation framework has been developed [Hallam, 1985] [Forster et al, 1988] to exploit variable type and quality of navigation data and methods. A self-navigation algorithm has been designed and tested in a simulated underwater environment notably making use of sonar derived environment data to infer motion with respect to the environment as well as simultaneously generating a world model. The navigator structure is flexible in that it can further accomodate *a priori* data, dead-reckoning data, navigation data obtained by beacons etc. The data is combined (or fused) by manipulating corresponding measurements and variances using Kalman filter arrangements. The system is robust in that it attempts to perform as well as possible with the data that is available to it.

A possible scheme for the incorporation of this self-navigation process into a general autonomous vehicle system is portrayed in Figure 8. The arrows connecting the boxes represent the general data flow, but the process of exploitation described earlier is less easy to represent. Algorithms generating route guidance strategies may need to exploit not just the collision avoidance process, but also control it by methods of sensor system 'suppression' and 'hallucination'. The simplicity of the collision avoidance behaviour may lead to difficulties in arranging fine motion control in confined locations. If the higher levels are confident about the nature of the environment, then suppression of the sensitivity of the collision avoidance process will lower resistance to motions to be achieved in confined space.

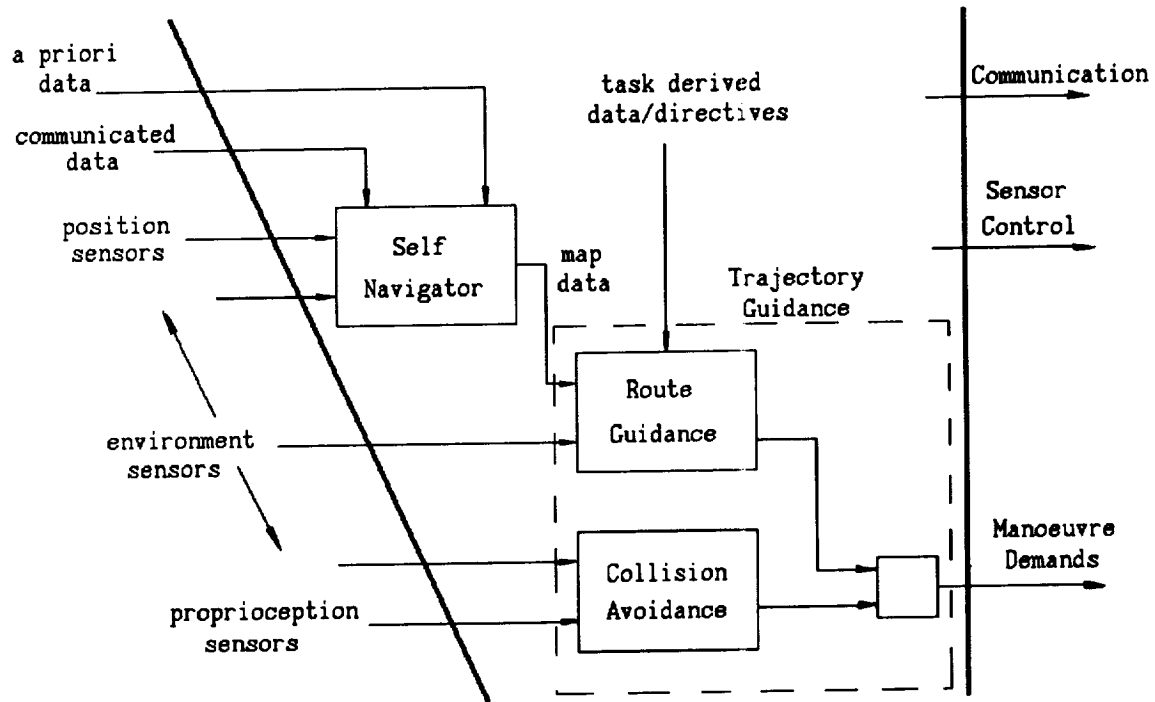


Figure 8: An architecture for mobile robot self-guidance

## 5 Conclusions

This paper has identified and put forward solutions to crucial issues regarding the competent performance of an autonomously operating robot, namely that of handling multiple and variable data sources containing overlapping information and maintaining coherent operation whilst responding adequately to changes in the environment. The validity of the ideas expressed have been supported by simple simulation experiments in the field of mobile robot navigation and guidance.

## 6 References

- [Brooks, 1986]  
Brooks, R, "Achieving Artificial Intelligence through building robots"  
M.I.T., A.I. Memo 899, May 1986;
- [Durrant-Whyte, 1987]  
Durrant-Whyte, H. F., "Uncertain geometry in robotics",  
Proc. IEEE Conference on Robotics and Automation, Vol. 2, p851, 1987.
- [Forster et al, 1988]  
Forster, P, Hallam, J, Smart, J and Howe, J,  
"Feature-based navigation for an underwater vehicle",  
Proceedings of the 2nd International Workshop on Subsea Robotics, Japan, 1988.
- [Hallam, 1985]  
Hallam, J C T, "Intelligent automatic interpretation of active marine sonar",  
PhD thesis, Edinburgh University, 1985.
- [Raibert, 1986]  
Raibert, M, "Legged robots that balance", M.I.T. Press, 1986.
- [Stillings et al, 1987]  
Stillings, N, Feinstein, M, Garfield, J, Rissland, E, Rosenbaum, D, Weisler, S & Baker-  
Ward, L, "Cognitive Science - An introduction", M.I.T. Press, 1987.